ZERO GYRO KALMAN FILTERING
IN THE PRESENCE OF A REACTION WHEEL FAILURE

Sun Hur-Diaz[*], John Wirzburger[†], Dan Smith[§], Mike Myslinski[£]

Typical implementation of Kalman filters for spacecraft attitude estimation involves the use of gyros for three-axis rate measurements. When there are less than three axes of information available, the accuracy of the Kalman filter depends highly on the accuracy of the dynamics model. This is particularly significant during the transient period when a reaction wheel with a high momentum fails, is taken off-line, and spins down. This paper looks at how a reaction wheel failure can affect the zero-gyro Kalman filter performance for the Hubble Space Telescope and what steps are taken to minimize its impact.

INTRODUCTION

The Hubble Space Telescope (HST) was launched back in 1990, and with four space shuttle servicing missions, it continues to operate and provide science observations. The third planned servicing mission was actually split into two shorter missions in order to accelerate the replacement of failed gyros during SM-3A. The last servicing mission, SM-3B, was conducted in 2002. Since then, HST has experienced failures of two of its six gyros. With the uncertainty of the next servicing mission, a concerted effort to extend the life of HST was initiated leading to the development of the two-gyro science (TGS) mode which offered the same science pointing performance of less than 7 milliarcsec jitter. In August 2006, the TGS mode became the operational mode with two of the remaining gyros turned off in order to extend their life. The development and operation of TGS is described in Reference 1.

Another product of the life extension initiatives was the Universal Kalman Filter (UKF) in preparation for a one-gyro science (OGS) mode. The universal nature of the Kalman filter refers to its use in both science mode and safe mode. UKF has the capability to provide attitude and rate estimates without gyros. It processes measurements from magnetometers, coarse sun sensors, and any number of gyros, including zero. In the absence of three-axis rate information from the functional gyros, a high fidelity dynamics model had to be utilized to achieve the Kalman filter (KF) performance required for both science and safe mode operations.

[*] Emergent Space Technologies, Inc., Greenbelt, MD 20770. E-mail : sun.hur-diaz@emergentspace.com.

[†] Honeywell Technology Solutions, Inc., Columbia, MD 21044. E-mail: jwirzburger@hst.nasa.gov.

[§] Lockheed Martin Mission Services, Greenbelt, MD 20770. Email: dsmith@hst.nasa.gov.

[£] Honeywell Technology Solutions, Inc., Columbia, MD 21044. E-mail: mmyslinski@hst.nasa.gov.

Three-axis HST control actuation is nominally achieved through a set of four reaction wheels which provide limited redundancy. Three of the wheels have been in near constant operation since launch, over 150,000 hours. One has been replaced twice, first in SM-2 and second in SM-3B. Given that the end of HST operations is projected to be 2014, there is a possibility that a wheel failure could occur during the science mission. The operation of HST for science observations with less than three wheels is currently under study.
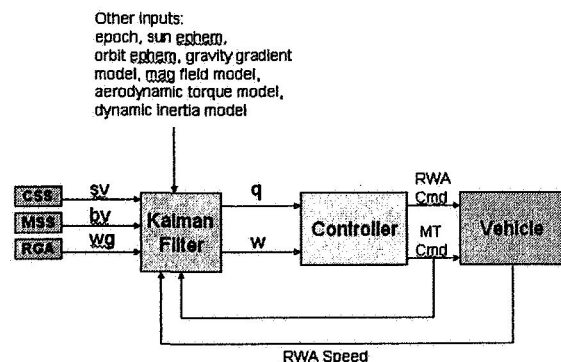
The implications of a wheel failure to the Kalman filter could be filter divergence unless appropriately accounted for. In normal operations, each individual wheel tachometer provides accurate speed information which the Kalman filter processes to determine the total system momentum as well as the torque provided by the reaction wheels. This is used in the dynamics model for the time update computation of the Kalman filter. In the event a reaction wheel fails and is powered off due to a safe mode test response, the tachometer output of the unpowered wheel will be erroneous. For a reaction wheel failure at high speeds, the resulting error in momentum calculation can cause large errors in the Kalman filter position and rate estimates. The filter is especially susceptible during orbit nights when only the magnetometer measurement is available.

This paper describes the Kalman filter algorithm onboard the HST, the possible effects of a failed wheel in its ability to provide accurate estimates, and steps that were taken to mitigate its effects.

UNIVERSAL KALMAN FILTER DESCRIPTION

A high-level block diagram of the KF in closed loop is shown in Figure 1. The controller acts based on the quaternion and rate estimates from the KF. Measurements processed by the KF are any available from the Magnetic Sensing System (MSS), the Coarse Sun Sensors (CSS), and the Rate Gyro Assembly (RGA). The actuators available are the Reaction Wheel Assembly (RWA) and the Magnetic Torquing System (MTS).

The design of a Kalman filter is discussed much in the literature and hence is only summarily described. A good reference is by Gelb (Ref. 2). The basic design used on HST is that of the Multiplicative Extended Kalman Filter (MEKF) described in Reference 3. The term "multiplicative" refers to the method in which the attitude quaternion is updated through a quaternion multiplication of an attitude error

Figure 1 Kalman filter in the loop

quaternion and the previous quaternion estimate. It differs from the "additive" method in

which the attitude error quaternion is added to the previous estimate to compute the updated quaternion.

In MEKF, the unity constraint of the attitude quaternion is maintained by estimating a 3-component attitude error, $\vec{a}$, about a reference unit quaternion, $\vec{q}_{ref}$. So the true quaternion can be represented as

$$\vec{q} = \vec{q}_e(\vec{a}) \otimes \vec{q}_{ref} \tag{1}$$

where $\vec{q}_e$ is the error quaternion approximated by

$$\vec{q}_e \approx \begin{bmatrix} \vec{a}/2 \\ 1 \end{bmatrix}$$
$$\vec{q}_e = \vec{q}_e / |\vec{q}_e| \tag{2}$$

and $\otimes$ represents quaternion multiplication. Note that the quaternion convention followed in this paper is that the first three elements represent the vector component and the fourth element is the scalar component.

The time derivative of $\vec{a}$ is given by

$$\dot{\vec{a}} = \vec{\omega} - \vec{\omega}_{ref} - \tfrac{1}{2}(\vec{\omega} + \vec{\omega}_{ref}) \times \vec{a} \tag{3}$$

where $\vec{\omega}_{ref}$ corresponds to the reference quaternion. Note that $\vec{q}_{ref}$ and $\vec{\omega}_{ref}$ are constantly updated through a reset procedure which is described later in this paper.

Basic Kalman Filter Equations

The Kalman filter consists of the measurement update and the time update. The measurement update consists of computing the Kalman filter gain matrix and updating the state estimates and their error covariance matrix. The gain matrix is given by

$$K_k = P_k(-)Y_k^T \left[Y_k P_k(-)Y_k^T + R_k \right]^{-1} \tag{4}$$

where $P_k(-)$ is the covariance matrix propagated to the measurement epoch in the previous time update, Y_k is the measurement sensitivity matrix evaluated using the states propagated to the measurement epoch,

$$Y_k \equiv \frac{\partial \vec{y}(\vec{x}, t_k)}{\partial \vec{x}} \bigg|_{\vec{x} = \vec{x}_k(-)} \tag{5}$$

3

and R_k is the measurement noise covariance. The gain matrix is used to update the state estimate as follows:

$$\vec{x}_k(+) = \vec{x}_k(-) + K_k[\vec{z}_k - \vec{y}(\vec{x}_k(-), t_k)] \tag{6}$$

where $\vec{z}_k$ is the measurement and is modeled as a nonlinear function of the predicted states at the time of the measurement plus noise:

$$\vec{z}_k = \vec{y}(\vec{x}_k(-), t_k) + \vec{v}_k, \quad \vec{v}_k \sim N(0, R_k) \tag{7}$$

where noise $\vec{v}$ is assumed to have normal (Gaussian) distribution with zero mean and covariance R, and $\vec{y}_k$ is the predicted measurement based on the propagated states.

The difference given in the square brackets of Eq. (6) is called the measurement innovations. In steady-state, the innovations are expected to be within some range. This expectation is the basis for the innovations check to eliminate measurements with grossly large errors before they are used in the Kalman filter and corrupt its estimates.

The error covariance associated with the measurement update is also computed as follows:

$$P_k(+) = [I - K_k Y_k]P_k(-)[I - K_k Y_k]^T + K_k R_k K_k^T \tag{8}$$

based on Joseph's algorithm to maintain symmetry and positive definiteness of the covariance matrix.

The time update of the Kalman filter involves propagating the states and the error covariance matrix to the next measurement epoch. The system is modeled by the time derivative of its states, and their noise $\vec{w}$ is assumed to have a normal (Gaussian) distribution with zero mean and spectral density Q:

$$\dot{\vec{x}} = \vec{f}(\vec{x}, t) + \vec{w}(t), \quad \vec{w}(t) \sim N(0, Q(t)) \tag{9}$$

The time update of the states is regarded as a continuous operation, and the system equation is propagated using a numerical integration. The time update of the error covariance matrix is performed as follows

$$P_{k+1}(-) = \Phi_k P_k(+) \Phi_k^T + Q_k \tag{10}$$

where Φ_k is the state transition matrix relating the states at time $k+1$ to the states at time k:

$$\bar{x}_{k+1} = \Phi_k \bar{x}_k \qquad (11)$$

and Q_k is the covariance matrix associated with the spectral density $Q(t)$.

Measurement Equations

The measurements considered in the UKF consist of rate measurements from the RGA and the unit magnetic field vector and the unit sun vector both in the vehicle frame as derived from the MSS and the illuminated CSS, respectively. The CSS measurements are processed only when the sun is within the linear range of roughly 41 deg half cone angle to ensure accurate sun vector measurements. All measurements are sampled at 1 Hz.

The measurement equation for the unit magnetic field vector is:

$$\begin{aligned} \bar{y}_{MSS} = \hat{b}_V &= A(\bar{q})\hat{b}_I + \bar{v}_{MSS} \\ &= A(\bar{q}_e(\bar{a}))A(\bar{q}_{ref})\hat{b}_I + \bar{v}_{MSS} \end{aligned} \qquad (12)$$

where $\bar{q}_e$ is given by Eq. (2) and

$\hat{b}_I$ = unit mag field vector in inertial frame computed from model, epoch, and orbit

$A(\bar{q})$ = direction cosine matrix (DCM) from inertial to vehicle frame

$$= \begin{bmatrix} q_1^2 - q_2^2 - q_3^2 + q_4^2 & 2(q_1q_2 + q_3q_4) & 2(q_1q_3 - q_2q_4) \\ 2(q_1q_2 - q_3q_4) & -q_1^2 + q_2^2 - q_3^2 + q_4^2 & 2(q_2q_3 + q_1q_4) \\ 2(q_1q_3 + q_2q_4) & 2(q_2q_3 - q_1q_4) & -q_1^2 - q_2^2 + q_3^2 + q_4^2 \end{bmatrix}$$

$\hat{b}_V$ = unit mag field vector in vehicle frame

$\bar{v}_{MSS}$ = 3×1 measurement noise for the unit mag field vector measurement

The measurement equation for the sun vector is similar except that an extra rotation due to measurement biases is included. The predicted CSS measurement is given by

$$\begin{aligned} \bar{y}_{CSS} = \hat{s}_V &= A(\bar{q}_b)A(\bar{q})\hat{s}_I + \bar{v}_{CSS} \\ &= A(\bar{q}_b(\bar{a}_b))A(\bar{q}_e(\bar{a}))A(\bar{q}_{ref})\hat{s}_I + \bar{v}_{CSS} \end{aligned} \qquad (13)$$

where

$$\hat{s}_I = \text{unit sun vector in inertial frame obtained from sun ephemeris}$$

$$A(\vec{q}_b) = \text{DCM corresponding to sensor biases in vehicle frame}$$

$$\hat{s}_V = \text{unit sun vector in vehicle frame}$$

$$\vec{v}_{CSS} = 3 \times 1 \text{ measurement noise for the unit sun vector measurement}$$

The measurement bias is to account for low-frequency errors due to uncalibrated errors like albedo, and is modeled as a first-order Gauss Markov process as follows:

$$\dot{\vec{a}}_b = -\vec{a}_b / \tau_b + \vec{w}_{ab} \tag{14}$$

where τ_b is the correlation time constant of the measurement bias state, and $\vec{w}_{ab}$ is the noise representing how well the model approximates the low-frequency errors. The quaternion associated with the measurement bias state is approximated by

$$\vec{q}_b \approx \begin{bmatrix} \vec{a}_b / 2 \\ 1 \end{bmatrix}$$

$$\vec{q}_b = \vec{q}_b / |\vec{q}_b| \tag{15}$$

The gyro measurement model is straightforward. As many gyro measurements as available can be processed. The gyro measurement equation is given by

$$\vec{y}_{GYRO} = \begin{bmatrix} \omega_{GYRO\_1} \\ \omega_{GYRO\_1} \\ \vdots \\ \omega_{GYRO\_N} \end{bmatrix} = \begin{bmatrix} T_{Veh\_to\_GYRO\_1} \\ T_{Veh\_to\_GYRO\_2} \\ \vdots \\ T_{Veh\_to\_GYRO\_N} \end{bmatrix} \vec{\omega}(-) + \begin{bmatrix} \omega_{Drift\_1} \\ \omega_{Drift\_2} \\ \vdots \\ \omega_{Drift\_N} \end{bmatrix} + \vec{v}_{GYRO} \tag{16}$$

where $T_{Veh\_to\_GYRO}$ is the transformation from the vehicle frame to the gyro axis, ω_{Drift} is the gyro drift which is available from a separate drift calibration, and $\vec{v}_{GYRO}$ is the gyro noise.

Time Update Equations

The states of the Kalman filter are the vehicle inertial quaternion, the 3-axis rate in the vehicle-frame, and the three sun vector measurement bias states. The MEKF consists of the following nine states:

$$\bar{x} = \begin{bmatrix} \vec{a} \\ \vec{\omega} \\ \vec{a}_b \end{bmatrix} \qquad (17)$$

whose time derivative is given by

$$\bar{f} \equiv \begin{bmatrix} \vec{f}_1 \\ \vec{f}_2 \\ \vec{f}_3 \end{bmatrix} = \begin{bmatrix} \vec{\omega} - \vec{\omega}_{ref} - \frac{1}{2}(\vec{\omega} + \vec{\omega}_{ref}) \times \vec{a} \\ I^{-1}\left(-S(\vec{\omega})I\vec{\omega} + S(\vec{h})\vec{\omega} + \vec{G}(\vec{a}) - \dot{\vec{h}} + \vec{T}_C + \vec{T}_{AERO}\right) \\ -\vec{a}_b / \tau_b \end{bmatrix} \qquad (18)$$

where I is the inertia matrix which can be time-varying due to the solar array position changes, $\vec{h}$ is the RWA angular momentum, $\vec{T}_c$ is the commanded torque to the magnetic torquers computed from the magnetic moment commands to the torquer bars and the MSS-measured magnetic field, $\vec{T}_{AERO}$ is the expected aerodynamic torque computed from a model with vehicle velocity vector and solar array position, and $\vec{G}$ is the gravity gradient term given by

$$\vec{G}(\vec{a}) = \frac{3\mu}{R_o^3} S(\hat{r}_v(\vec{a})) I \hat{r}_v(\vec{a})$$

$$\hat{r}_v(\vec{a}) = A(\vec{q}_e(\vec{a})) A(\vec{q}_{ref}) \hat{r}_I \qquad (19)$$

$$\approx \left[I_{3\times3} - S(\vec{a}) \right] A(\vec{q}_{ref}) \hat{r}_I \quad \text{(to first - order in } \vec{a}\text{)}$$

where R_o is the magnitude of the orbit radius, $\hat{r}_I$ is the unit radius vector in inertial frame, and μ is the gravitational constant. The symbol $S(\ )$ represents the cross-product matrix of a vector:

$$S(\vec{v}) \equiv \begin{bmatrix} 0 & -v_3 & v_2 \\ v_3 & 0 & -v_1 \\ -v_2 & v_1 & 0 \end{bmatrix} \qquad (20)$$

In addition to the nine states, the reference states are propagated by the following equations:

$$\dot{\vec{q}}_{ref} = \tfrac{1}{2}\begin{bmatrix} \vec{\omega}_{ref} \\ 0 \end{bmatrix} \otimes \vec{q}_{ref}$$

$$\dot{\vec{\omega}}_{ref} = I^{-1}\left(-S(\vec{\omega}_{ref})I\vec{\omega}_{ref} + S(\vec{h})\vec{\omega}_{ref} + \vec{G}(\vec{a}) - \dot{\vec{h}} + \vec{T}_C + \vec{T}_{AERO}\right) \qquad (21)$$

The wheel angular momentum, $\vec{h}$, is computed from the measured wheel speeds multiplied by the wheel inertias and then converted to the vehicle frame. The time rate of change of the wheel angular momentum is computed from the back-difference of this wheel angular momentum.

Reset Operation

In the MEKF, the following reset operation is performed as part of the measurement update after $\vec{a}_k(+)$ and $\vec{\omega}_k(+)$ have been computed:

$$
\begin{aligned}
&1. \quad \vec{q}_{ref}(+) = \vec{q}_e(\vec{a}_k(+)) \otimes \vec{q}_{ref}(-) \\
&2. \quad \vec{q}_{ref}(+) = \vec{q}_{ref}(+)/\left|\vec{q}_{ref}(+)\right| \\
&3. \quad \vec{\omega}_{ref}(+) = \vec{\omega}_k(+) \\
&4. \quad \vec{a}_k(+) = 0
\end{aligned}
\qquad (22)
$$

Note that the reset operation is essentially a measurement update of the reference quaternion and the reference angular velocity after which the attitude error state is reset to zero.

## Convergence and Innovations Check

Although the likelihood of getting very bad measurements from the MSS and the CSS is very small, spikes in measurements from glints or day/night transitions should be eliminated and not processed by the Kalman filter if possible. Monitoring of the estimate error covariance and measurement innovations and residuals can enable detection of such errors that are outside of the expected range.

Impact of a Failed Wheel

During normal operations, RWA speeds are read directly from the tachometer signals from each RWA into Flight Software (FSW). In event a RWA fails and is powered off due to Safing test response, the tachometer output of the unpowered RWA will be erroneous. Analysis of UKF performance in simulations that include an RWA failure at high speed indicates that the resulting errors in momentum calculation cause large errors in the UKF rate and position estimates.

Figure 2 shows the simulation result of a case where the zero-gyro Kalman filter diverges with the initial speed of the failed wheel around 2500 rpm. The top plot shows the magnitude of the position estimate error along with the orbit night flag. The bottom plot shows the error in the knowledge of the failed wheel's speed. Note that the position estimate error decreases as the wheel speed error decreases. Although the Kalman filter recovers, controlling during the high-error period will produce unacceptable results.
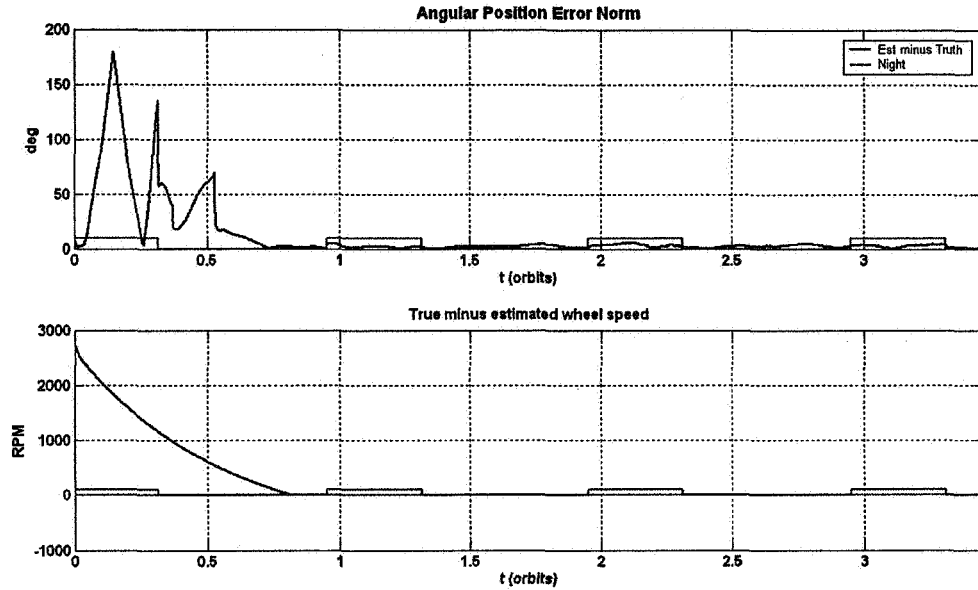
**Figure 2 An example of Kalman Filter divergence due to wheel failure**

MEDIATION TECHNIQUES

To reduce these errors, a four-fold approach was taken. First, a model of the wheel deceleration after power off was developed based on data provided by the reaction wheel manufacturer. Second, in order to minimize errors in this model to capture the conditions at the start of the deceleration, an accurate model of the wheel torque output was developed. Third, a more sensitive safe mode test was designed to catch any failures early and shorten the response time. The fourth step involves switching of certain KF parameters to account for reduced accuracy in the dynamics model. These four mitigation steps improved the performance of the Kalman filter during the transient spin-down period of a failed wheel.

RWA Spin Down Estimate

When power is removed from the RWA, torque on the wheel shaft provided by the RWA motor is removed leaving only mechanical drag torque and, at wheel speeds greater than approximately 2000 rpm, electro-magnetic drag to act upon the wheel shaft. These torques oppose the wheel shaft motion and cause the wheel to spin down. When the RWA is separated from vehicle bus power while spinning, it becomes an electrical power

9

generator, supplying power to internal circuits that act as a load and drag on the wheel shaft. At a certain rotation rate, the power supplied from the wheel is insufficient to sustain the internal circuits, and under-voltage protection in the RWA circuits will remove the electrical load from the RWA shaft. The wheel speed at which this occurs will vary between RWAs due to the tolerances in individual components that make up the analog circuits of the RWA, however the range of speeds at which the electrical load will be removed is estimated to be from 1500 to 2500 rpm. Mechanical load from the bearings may also vary between RWAs. Figure 3 shows the vendor provided deceleration envelope based on the range of drag torque variation between RWAs.
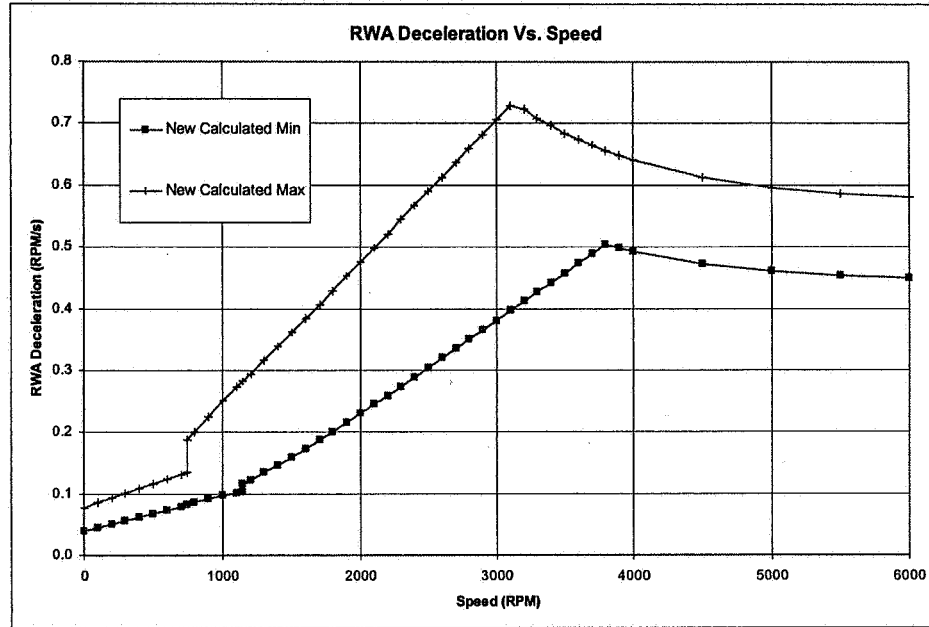

Figure 3 Vendor supplied RWA Spin Down Deceleration Envelope

Equations for the RWA spin down (truth) envelope and RWA speed estimate models were derived from curve fits of the deceleration envelope data. These equations were developed into a simulation for evaluation and were found to be adequate representations of the vendor supplied data. Figures 4 and 5 are plots of the calculated RWA accelerations and RWA speeds respectively for the hardware envelope and estimate.

10
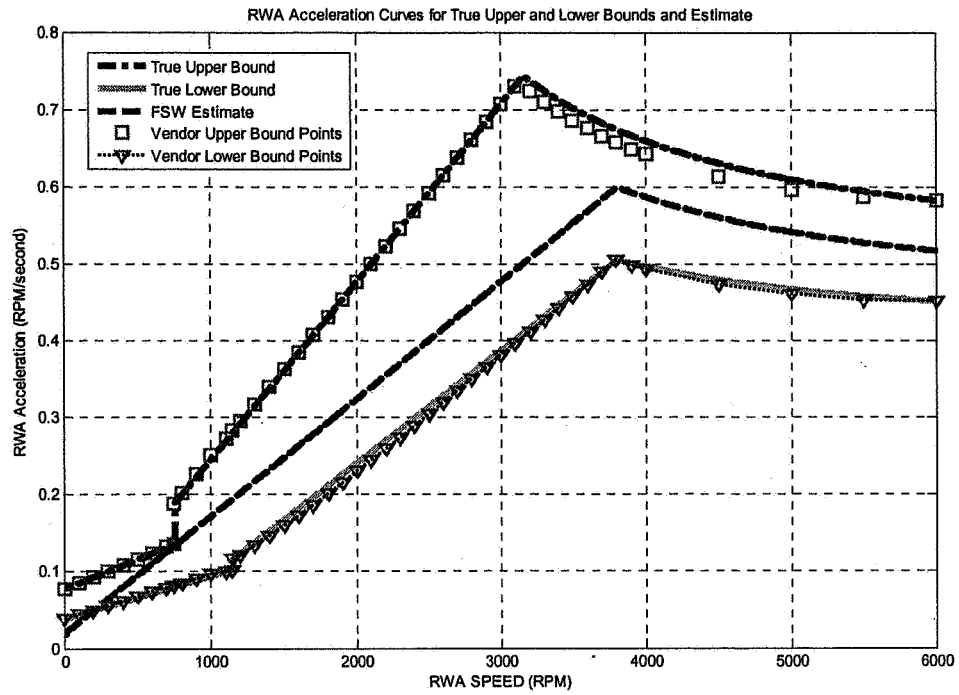
Figure 4 RWA Spin Down Deceleration
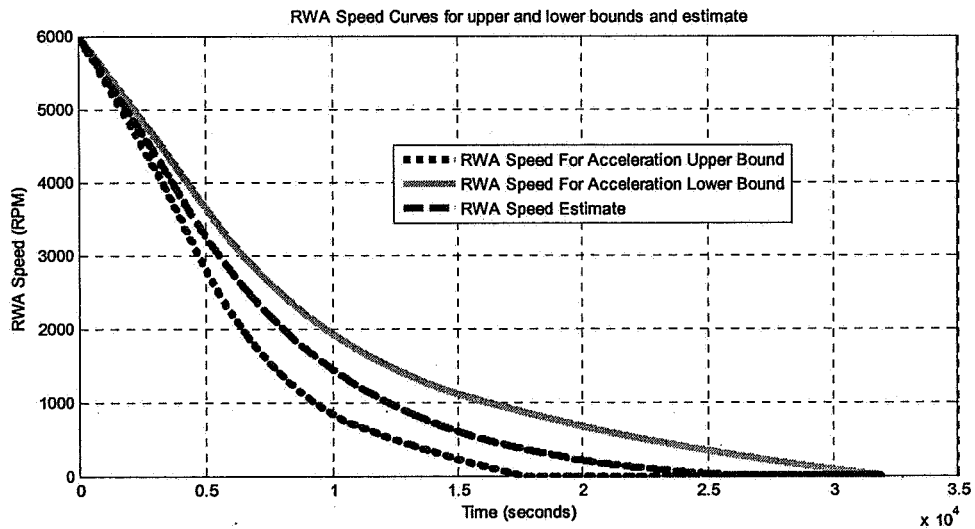
Figure 5 RWA Speed During Spin Down

The spin down characteristics were modeled in two segments, with a non-linear equation for the high speed segment and a line equation for the low speed segment.

11

RWA Torque Estimator

The change in RWA speed, and thus momentum, is effected by FSW torque commands to the RWA. Analog circuitry within the RWA integrates the 40 Hz torque commands to provide a targeted wheel speed to an internal speed control circuit, with feedback provided by a wheel mounted tachometer. RWA torque output will generally match the commanded torque, which is limited to about 0.82 Newton-meters (N-m), when RWA speed magnitude is less than approximately 3300 RPM. Due to the tolerances of components within the analog circuitry of the RWA, there is some variability in the linear range of the torque output.

Above 3300 rpm the wheel speed control circuit output saturates in the direction that opposes acceleration reducing the total torque command. The current loop gain decreases with speed causing an offset that opposes acceleration that is a function of speed. Mechanical drag of about .03 N-m will no longer be compensated for by the speed control circuit and thus will also affect the resultant wheel output torque.

At higher speeds, back-EMF within the RWA electronics results in decreasing torques with increasing speed until a maximum RWA speed magnitude is reached, where the output torque level goes to zero. The point at which the back-EMF start affecting the torque output of the wheel is a function of the bus voltage being supplied to the RWA. At 32 volts the back-EMF kicks in at about 4800 rpm and limits the wheel speed to 6600 rpm. For a 24 volt bus, back-EMF starts at about 3600 rpm and results in a zero output torque at about 5100 rpm.

Figure 6 shows the output torque at positive wheel speeds for +0.82, +0.41, 0.0, -0.41, and -0.82 N-m torque commands. The solid lines on the graphs were generated by the vendor using a PSpice model of the reaction wheel. The dashed graph lines are curve fits to the PSpice data which are a function of the commanded torque, current wheel speed, and bus voltage. This empirical fit of the data has a maximum deviation from the PSpice output of less than .04 N-m.

12
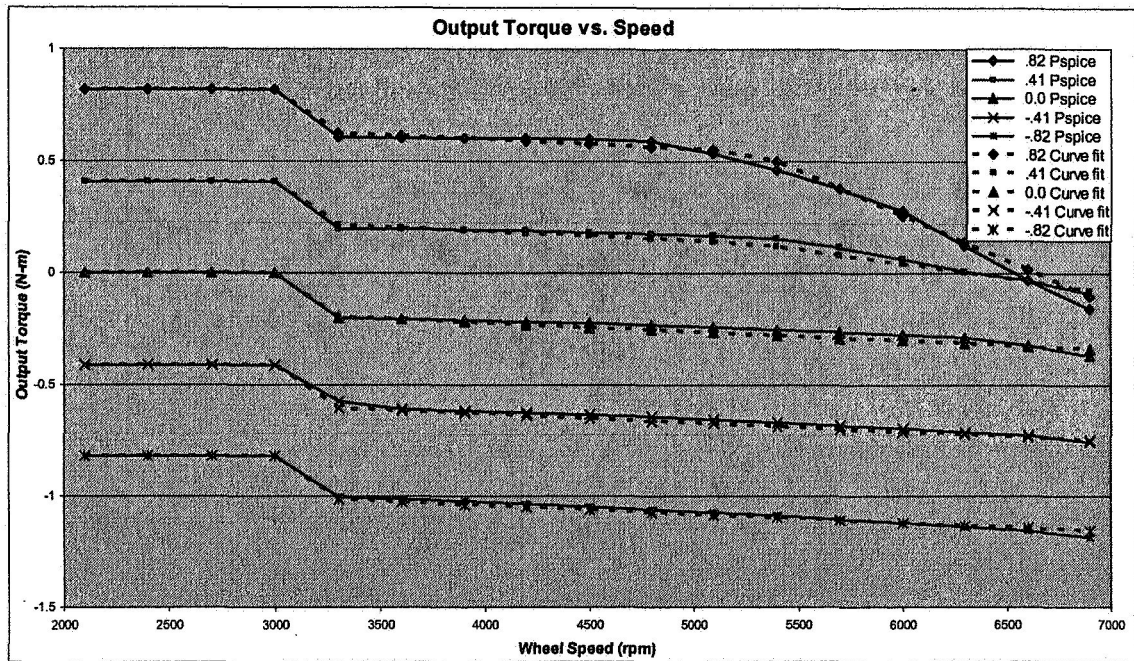
**Figure 6 Output torque as a function of commanded torque**

Figure 7 shows how the torque output varies with the voltage supplied to the wheel motor. The point at which the output torque falls below zero for a positive torque command is the maximum speed the wheel can be commanded to. As can be seen on the plots, this saturation speed is lower for lower voltages.
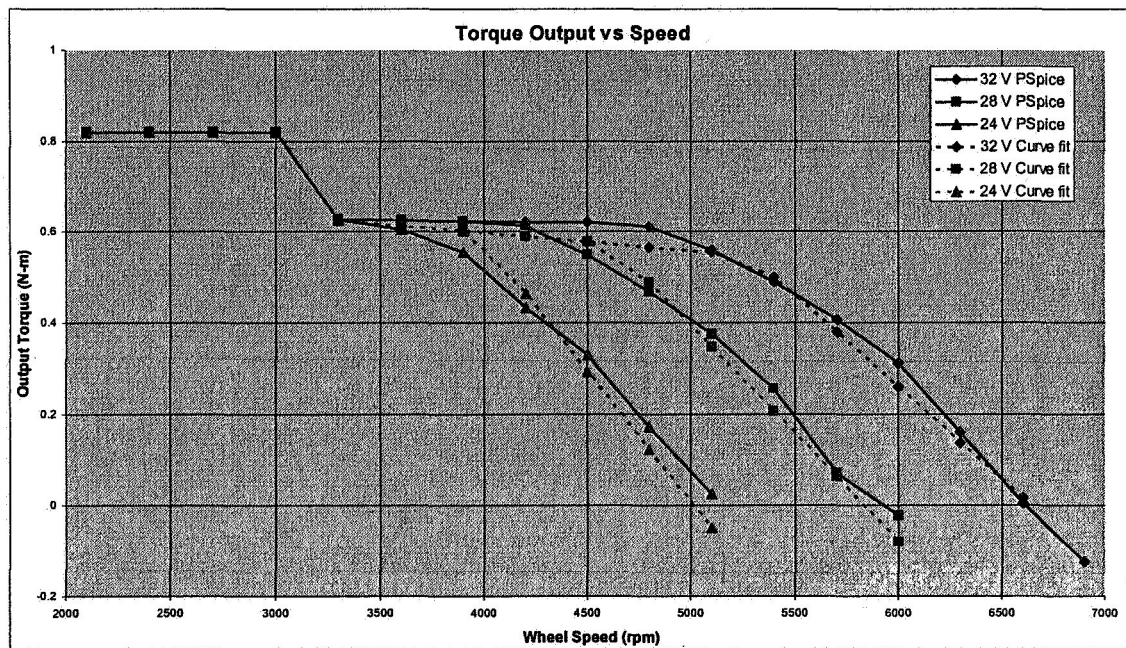
Figure 7 Output torque as a function of supply voltage

13

While there is good agreement between the curve fits and the PSpice output, this accuracy comes at the expense of software complexity, requiring calculation to determine where the back-EMF effect starts and the end-point of the torque roll-off, both of which are functions of the bus voltage. Further complicating the algorithms is the fact that two of the reaction wheels are nominally connected to a different power bus from the other two which can be at different voltages. It was therefore decided to simplify the output torque equations for the FSW implementation, and accept the larger errors between the estimator output and the actual reaction wheel speeds. Figures 8 and 9 compare the Torque Output plots produced by the curve fits with those from the simplified equations.

The slope and end-point values for the back-EMF range of the Estimator curve were selected to minimize the errors between the curve fits. Good agreement is seen between the original data and the simplified equations in the linear range and the speeds before the back-EMF attenuation range, but because of the strong relationship of the back-EMF effect and the voltage, much larger differences between the simplified equations and the vendor data occur when the variability of the voltage is not taken into account.
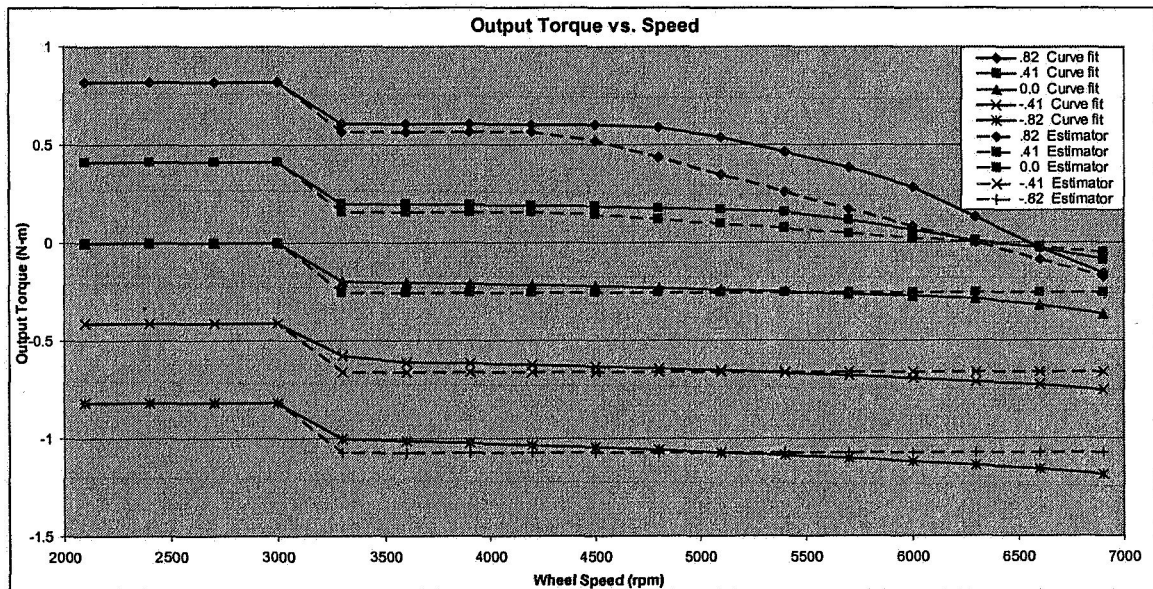
Figure 8 Estimator output compared to curve fit data

This error was minimized by calculating a line segment that splits the difference between it and the original curve fits for the lowest and highest voltages the wheels will see when commanded in the HST primary control modes. The maximum difference of .23 N-m between the original curve fits and the simplified equations occurs around 5400 rpm.

14
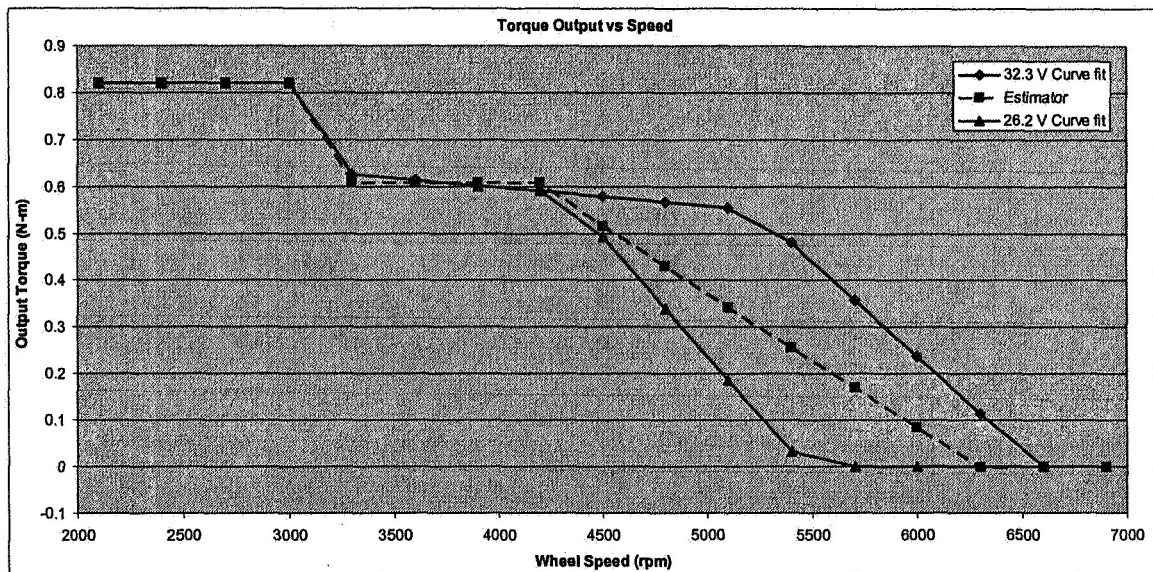
Figure 9 Estimator output compared to curve fit data

The torque computed from this model is incorporated into the RWA momentum estimator to accurately determine the wheels speed.

RWA Safing Test

One of the primary safing tests on the RWA is the RWA Momentum test which was modified to take advantage of the more accurate RWA torque calculation to make the test more robust to detecting RWA anomalies in a timely manner.

The first part of the test, which runs every 10 seconds, calculates the difference between the estimated speed and the actual tachometer output and compares it to the appropriate limit. If the check passes then the failure counter is cleared and the estimated value of the RWA momentum is set equal to the tachometer value in order to keep the estimator synchronized with the actual wheel momentum. If the check fails, a failure counter is incremented and an event message is placed in vehicle telemetry. The next time through the flow the limits are again determined, but this time the limit is set to twice the data base value. This is done because the limit check is being done against 20 seconds of integrated torque commands, so the limits are set to twice the value used for the 10-second integration. If the test fails a second time then the failure counter is incremented, a warning message is issued, and the limit is increased appropriately. The process is repeated every 10 seconds until an acceptable momentum is calculated and the counter is cleared or until the second part of the RWA Momentum test fails.

The second part of the RWA Momentum test checks the value of the failure counter, and if it exceeds a database defined limit, will cause the test to fail and result in the reaction wheel to be taken out of the control loop and turned off.

15

The momentum limit for linear range of the reaction wheel is based on observed RWA telemetry noise and errors between commanded torque and wheel output torque. The limit for the non-linear range is set a conservative .2 N-m-s higher to account for uncertainties in the RWA behavior at the higher momentums. A circuit in the wheel electronics that is not modeled in the equations is the synchronization circuit, which can cause the speed loop to subtract .17 N-m from the commanded torque level for close to 40 seconds after the actual wheel speed comes below the speed loop saturation point. Worst-case, this would result in the .17 N-m-s being applied until the speed is 260 rpm below the speed control loop limit. The safing test limit must therefore account for this possible error at the lowest speed where this could occur, 260 rpm below the speed control loop saturation speed.

## Kalman Filter Parameter Switching

When the RWA safing test fails and the failed wheel is taken out of the loop, some of the Kalman filter parameters are switched to different values to account for the higher uncertainty in the dynamics model. In particular, the process noise covariance Q is increased. Also, since the expected range of the measurement innovations and residuals is higher with the increase in estimate error, their limits for measurement rejection and filter convergence tests are increased during this period of higher uncertainty. When the failed wheel's speed decreases sufficiently, and its estimate goes below a minimum value, these Kalman filter parameters are set back to their nominal values.

RESULTS

The wheel deceleration model, the wheel torque model, the RWA momentum safing test, and the switching of the Kalman filter parameters were incorporated into a high-fidelity simulation. Monte Carlo evaluations of the models indicate that these mitigation steps allow the UKF to maintain an adequately accurate estimate of vehicle attitude and rate during RWA spin down.

The same case shown in Figure 2 was simulated first with only the models of the wheel momentum estimators during powered and unpowered states and the new safe mode test in place. The corresponding attitude estimate error and the wheel speed error are shown in Figure 10. The wheel speed error has dropped significantly and so has the Kalman filter position estimate error.
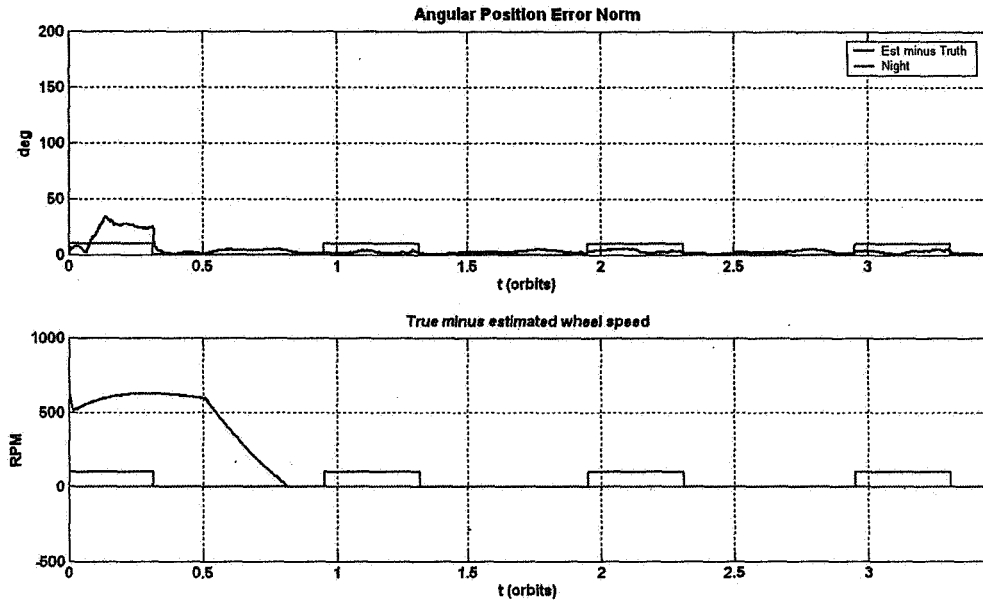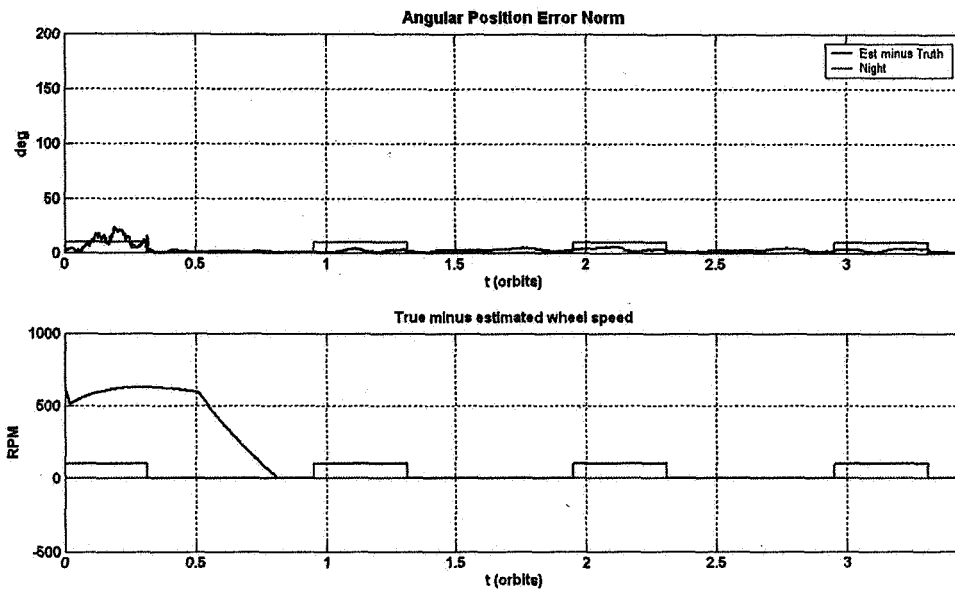
Figure 10 Kalman filter performance with wheel speed estimators and modified safing test

In spite of these models, however, the uncertainty on the wheel speed can still be several orders of magnitude higher than the measurement accuracy of the tachometer. The larger uncertainty is taken into account by increasing the Kalman filter process noise during the spin-down period. Then the nominal value of the process noise is reinstated once the estimate of the wheel speed reaches within a tolerance of zero. The filter contains measurement innovations checks as well as convergence checks based on the state covariance and measurement residuals. So as not to erroneously trip the convergence monitor limit or to reject good measurements, thresholds for these checks had to be increased as well.

Figure 11 show the position estimate error with the added switching logic in the process noise and the covariance, innovations, and residual thresholds where the process noise covariance was increased by roughly two orders of magnitude during the spin-down period. This provided additional improvement in the Kalman filter performance as shown in the first plot with no change in the wheel speed error.

17

Figure 11 Kalman filter performance with wheel speed estimators, modified safing test, and Kalman filter parameter switching

CONCLUSIONS

A reaction wheel failure can cause a gyro-less Kalman filter to diverge during the period when the failed wheel is spinning down and there is no input from the tachometer. Models of the reaction wheel during its powered and unpowered states as well as a modified safe mode test to capture the failure early were shown to improve performance by increasing the accuracy of the dynamics model. Additionally increasing the process noise and managing the innovations, residuals, and covariance checks during this period to account for the increased uncertainty in the dynamics model improved the Kalman filter performance even further.

REFERENCES

1. Prior, M. and L. Dunham, "System Design and Performance of the Two-Gyro Science Mode for the Hubble Space Telescope," presented at the 56[th] International Astronautical Conference, Fukuoka, Japan, October 17-21, 2005.
2. Gelb, A., *Applied Optimal Estimation*, The MIT Press, Cambridge, MA, 1974.
3. Markley, F. L., "Attitude Error Representations for Kalman Filtering," *Journal of Guidance, Control, and Dynamics*, Vol. 26, No. 2, March-April 2003, p. 311-317.